

The Inhomogeneous Ionizing Background Following Reionization

Andrei Mesinger^{1*} & Steven Furlanetto²

¹Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA

²Department of Physics and Astronomy, University of California, Los Angeles, CA 90095, USA

17 June 2009

ABSTRACT

We study the spatial fluctuations in the hydrogen ionizing background in the epoch following reionization ($z \sim 5$ –6). The rapid decrease with redshift in the photon mean free path (m.f.p.), combined with the clustering of increasingly rare ionizing sources, can result in a very inhomogenous ionizing background during this epoch. We systematically investigate the probability density functions (PDFs) and power spectra of ionizing flux, by varying several parameters such as the m.f.p., minimum halo mass capable of hosting stars, and halo duty cycle. In order to be versatile, we make use of analytic, semi-numeric and numeric approaches. Our models show that the ionizing background indeed has sizable fluctuations during this epoch, with the PDFs being a factor of few wide at half of the maximum likelihood. The clustering of sources dominates the width of the PDFs, so analytic models must take large-scale clustering into account. The distributions also show marked asymmetries, with a high-value tail set by clustering on small scales, and a shorter low-value tail which is set by the mean free path. The power spectrum of the ionizing background is much more sensitive to source properties than the PDF and can be well-understood analytically with a framework similar to the halo model (usually used to describe dark matter clustering). Nevertheless, we find that Ly α forest spectra are extremely insensitive to the details of the UVB, despite marked differences in the PDFs and power spectra of our various ionizing backgrounds. Assuming a uniform ionizing background only underestimates the value of the mean ionization rate, $\langle \Gamma_{12} \rangle$, inferred from the Ly α forest by a few percent. Instead, analysis of the Ly α forest is dominated by the uncertainties in the density field. Thus, our results justify the common assumption of a uniform ionizing background in Ly α forest analysis even during this epoch.

Key words: cosmology: theory – early Universe – galaxies: formation – high-redshift – evolution

1 INTRODUCTION

The metagalactic hydrogen-ionizing ultraviolet background (UVB) is probably the most influential, pervasive radiation field throughout most of cosmic history. Much investigation has gone into understanding its properties and its impact on the intergalactic medium (IGM) (see Meiksin 2007 for a recent review). At low redshifts ($z \lesssim 3$ –4), the ionizing photon mean free path (m.f.p.) is very large, and so each gas element effectively “sees” many sources of ionizing radiation. A uniform UVB seems to be an accurate approximation in this regime, with spatial fluctuations contributing only a few percent of the mean (Zuo 1992; Fardal & Shull 1993; Gnedin & Hamilton 2002; Meiksin & White 2004; Croft 2004; Bolton et al. 2006).

However, the m.f.p. decreases rapidly with redshift, with extrapolations predicting that it should approach values of $\lambda_{\text{mfp}} \sim 30$ –40 comoving Mpc by $z \sim 5$ –6 (e.g. Storrie-Lombardi et al. 1994; Miralda-Escudé 2003). Such modest values of the m.f.p. at

$z \gtrsim 5$ can induce significant fluctuations in the UVB, sourced by cosmic variance in the number of ionizing sources inside each λ_{mfp} horizon. Additionally, the $1/r^2$ flux attenuation, combined with the clustering of sources, can introduce additional scatter in the UVB on scales smaller than λ_{mfp} (Zuo 1992; Meiksin & White 2004). These fluctuations have been largely ignored in high-redshift analyses (though see §5 in Bolton & Haehnelt 2007 and §4 in Feng et al. 2008).

Here we undertake a systematic study of the post-reionization metagalactic ionizing background at $z \sim 5$ –6. We vary source prescriptions and λ_{mfp} in an attempt to *infer general conclusions about the theoretical properties and observational importance of spatial fluctuations in the UVB*. We focus on the simplest statistics, the probability density function (PDF) and power spectrum, which are most likely to be observed. Given the dearth of knowledge concerning the properties of absorbers and sources at high- z , we make use of a range of methods that include analytic, semi-numeric and numeric approaches in order to be comprehensive. We compare simple analytic models of these quantities to the numerical results in order to shed light on the origin and importance of the

* Hubble Fellow; email: mesinger@astro.princeton.edu

variations; we will show that analytic models are accurate in some, but not all, cases.

This paper is organized as follows. In §2, we outline the various modeling tools we use to study the fluctuations in the ionizing background. In §3 and §4 we present the PDFs and power spectra of various UVBs, respectively. In §5, we study the impact of a spatially varying UVB on the Ly α forest. Finally, in §6, we summarize our findings and offer conclusions.

We quote all quantities in comoving units, with the exception of flux, and we denote proper units with a prefix 'p'. We adopt the background cosmological parameters ($\Omega_\Lambda, \Omega_M, \Omega_b, n, \sigma_8, H_0$) = (0.72, 0.28, 0.046, 0.96, 0.82, 70 km s $^{-1}$ Mpc $^{-1}$), matching the five-year results of the *WMAP* satellite (Dunkley et al. 2008; Komatsu et al. 2008).

2 GENERATING INHOMOGENEOUS FLUX FIELDS

2.1 Semi-Numerical Simulations with Flux Tool

We generate our ionizing flux fields following the procedure described in Mesinger & Dijkstra (2008). We briefly outline the procedure below.

We begin with halo fields at $z = 5$ and $z = 6$ generated with the semi-numerical simulation DexM¹, which has been shown to reproduce the correct number density and clustering properties of halos well into the quasi-linear and non-linear regimes (Mesinger & Furlanetto 2007; Fig. 1 in Dijkstra et al. 2008; Mesinger et al. in preparation). Our simulation boxes are 150Mpc on a side, with a 150/1800 Mpc halo grid cell size. The velocity fields used to perturb the halo field were generated on a lower resolution 900 3 grid; thus our final halo field resolution is $\Delta x = 150/900 = 0.17$ Mpc.

For each halo field, we create a corresponding UV flux field on a 150 3 grid (spatial resolution of 1 Mpc), by performing a halo mass/ r^2 weighted sum. Specifically, we compute the flux of ionizing photons (in units of ionizing photons s $^{-1}$ pc m $^{-2}$) with

$$f(\mathbf{x}, z) = \frac{(1+z)^2}{4\pi} \sum_{M_i > M_{\min}} X_i(\text{DC}) \frac{M_i \times \epsilon_{\text{ion}}}{|\mathbf{x} - \mathbf{x}_i|^2} e^{-|\mathbf{x} - \mathbf{x}_i|/\lambda_{\text{mfp}}}, \quad (1)$$

where \mathbf{x} is the location of the cell of interest, M_i is the total halo mass, \mathbf{x}_i is its location, and the factor of $(1+z)^2$ converts the factor $|\mathbf{x} - \mathbf{x}_i|^2$ from comoving into proper units. We also include a duty cycle² parameter through the random variable $X_i(\text{DC})$, which has a value of 1 with likelihood DC or 0 with likelihood 1 – DC. Finally, the factor ϵ_{ion} in eq. (1) denotes the rate at which ionizing photons are released into the IGM by a dark matter halo per unit mass. Unless stated otherwise, we assume a fiducial value of $\epsilon_{\text{ion}}(z) = 3.8 \times 10^{58} / \text{DC} \left[\frac{\Omega_b}{\Omega_M} \frac{1}{t_H(z)} \right] \frac{\text{photons}}{M_\odot \text{s}}$, which provides a good fit to the observed luminosity functions of Ly α emitting galaxies (LAEs) (Shimasaku et al. 2006; Kashikawa et al. 2006; Dijkstra et al. 2007; Stark et al. 2007; McQuinn et al. 2007), and

¹ <http://www.astro.princeton.edu/~mesinger/DexM/>

² Note that there are several definitions of the term “duty cycle” in the literature, some of which include this same parameter in the star-formation timescale of each halo (an assumption only valid in extremely idealized scenarios). We adopt the simple definition above of the duty cycle, namely the fraction of dark matter halos which are “on” (i.e. hosting active galaxies emitting ionizing photons into the IGM; c.f. Lee et al. 2008).

the $z = 6$ Lyman Break galaxies (LBGs) (Bouwens et al. 2006). However, as we are mostly concerned with the shape of the flux PDF, where applicable we present results in units of the mean flux or intensity, $J/\langle J \rangle$.

In constructing the UVB, we also assume a minimum halo mass able to host stars, M_{\min} , and explore mainly two different values for this parameter: 1.6×10^8 and $1.4 \times 10^9 M_\odot$. The former value corresponds to a virial temperature of 10^4 K at these redshifts, and represents the regime of ineffective feedback on atomically cooled halos during reionization (Mesinger & Dijkstra 2008; Okamoto et al. 2008). The later value represents the regime of very inefficient star formation inside such small halos, perhaps due to strong radiative and/or mechanical feedback (e.g. Yepes et al. 1997; Scannapieco et al. 2006; Pawlik & Schaye 2008). The exact value of the higher M_{\min} was also motivated by the fact that the number density of halos with mass greater than $1.4 \times 10^9 M_\odot$ is roughly ten times smaller than the number density of halos with masses greater than $1.6 \times 10^8 M_\odot$ at $z \approx 5$, thus allowing us to compare results at fixed source number density with values of DC = 1.0 and 0.1 in the two models, respectively.

2.2 Hydrodynamical Simulations with Flux Tool

Additionally, we generate a flux field corresponding to the $z = 5.71$ source field from the cosmological hydrodynamical simulation presented in Trac et al. (2008) (their $z = 6$ reionization model). This simulation is fixed-grid, 143 Mpc on a side, and includes prescriptions for modeling dark matter, baryons and ionizing photons (for details see Trac & Cen 2007 and Trac et al. 2008). The density field was calculated on grid of 0.19 Mpc cells, which resolves the Jeans length in the mean density, ionized intergalactic medium (IGM) by a factor of few, and then smoothed to a cell size of 0.74 Mpc. Each halo’s instantaneous star-formation rate (SFR) is proportional to the instantaneous gas accretion rate, and enters our eq. (1) in place of the halo’s mass, M_i , with the normalization adjusted accordingly so as to match the mean value of the UVB. We make use of this flux field, combined with the $z = 5.71$ gas density fields from the hydro-simulation, to generate more accurate mean Ly α forest flux decrement statistics in §5. At these redshifts and scales, our semi-numerically generated density fields somewhat over-predict the rare voids which dominate the this statistic (Mesinger et al., in preparation). Furthermore, the ray-tracing algorithm from the numerical simulation over-predicts the fluctuations in the flux field, due to an insufficient number of ray splittings (Trac 2008, private communication). This problem only becomes severe following reionization, but since this is the epoch we study here, we use this hybrid prescription in some of the results below (i.e. we do not use the radiative transfer field from the simulation).

In Fig. 1, we present a 0.74 Mpc thick slice through the density and flux fields based on this $z = 5.71$ simulation output from Trac et al. (2008). The flux fields were calculated assuming $\lambda_{\text{mfp}} = 30$ Mpc, and DC = 1. Even for such a moderately high choice of λ_{mfp} , we can qualitatively see that there is significant inhomogeneity in the flux fields. Furthermore, it is evident that the flux and density fields are highly correlated, an issue we will return to in §5.

2.3 Analytic Model

We complement these numerical techniques with a relatively simple analytic model. While this approach cannot incorporate all of

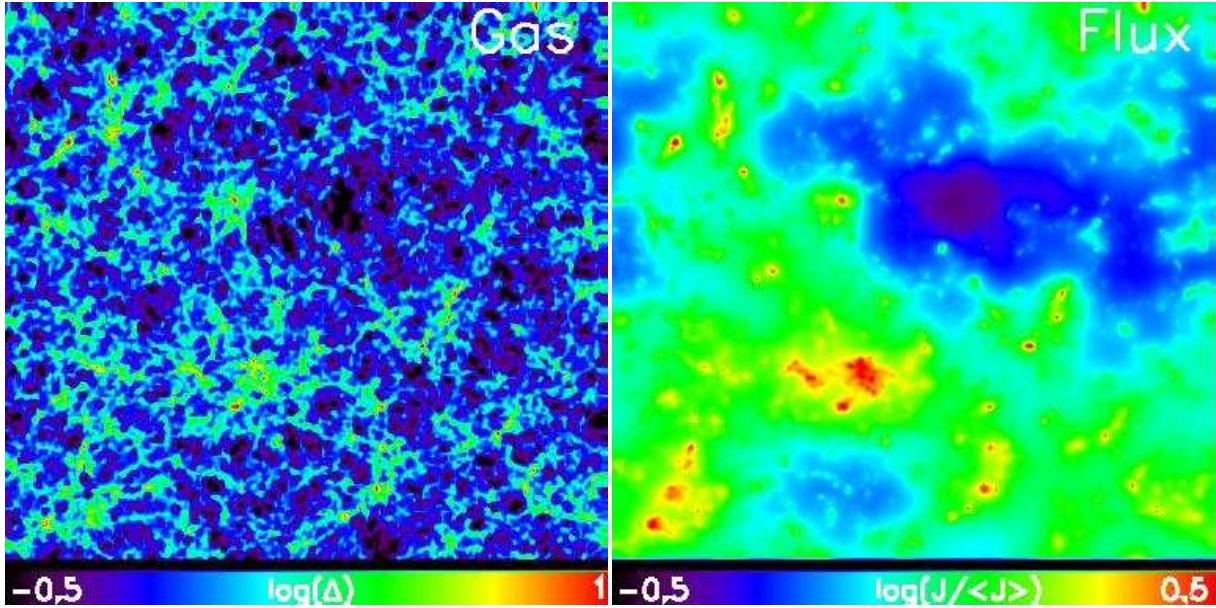


Figure 1. Slices through the $z = 5.71$ gas density (left) and UVB (right) fields using the hydrodynamical simulation of Trac et al. (2008) and our simple flux generation tool. Each slice is 143 Mpc on a side and 0.74 Mpc thick. The UV flux is generated using the star-formation rates obtained from the simulation and the prescription outlined in §2, assuming $\lambda_{\text{mfp}} = 30$ Mpc, DC=1, and $M_{\text{min}} \sim 3 \times 10^8 M_\odot$.

the physics provided by the semi-numeric and numeric models, we will find it to be helpful in elucidating the physics of the background radiation field. As in the remainder of this paper, we will focus on computing two statistical descriptions of the flux field: the PDF and the power spectrum.

2.3.1 Flux PDF

The PDF of the radiation background J , normalized to its mean value $\langle J \rangle$, can be computed exactly for randomly distributed sources, provided that we assume a constant attenuation length λ_{mfp} for ionizing photons much smaller than the Hubble length (Zuo 1992; Meiksin & White 2004). In this case, the crucial parameter is the number of sources within each attenuation length, $\bar{N}_0 = (4\pi/3)n_i\lambda_{\text{mfp}}^3$, where n_i is the source number density, and

$$\frac{dp(>X|\bar{N}_0)}{dX} = \frac{1}{\pi} \int_0^\infty ds \exp \left[-s\bar{N}_0 \int dx x\phi(x) \text{Im } G(sx) \right] \times \cos \left[-sX + s\bar{N}_0 \int dx x\phi(x) \text{Re } G(sx) \right] \quad (2)$$

where $X = J/J_*$, $\phi(x)$ is the source luminosity function, in units of the mean luminosity L_* and normalized to have unit integral, $J_* = L_*/(4\pi\lambda_{\text{mfp}})^2$,

$$G(t) = \int_0^\infty du \tau^3(u) e^{itu}, \quad (3)$$

and $u = e^{-\tau}/\tau^2$. In this limit, the mean background is $\langle X \rangle = 3\bar{N}_0$, or $\langle J \rangle \propto \bar{N}_0 J_* \propto L_* r_0$.

At the end of reionization, $\lambda_{\text{mfp}} \gtrsim 20$ Mpc and (typically) $\bar{N}_0 \gtrsim 10^5$. Thus, the Poisson fluctuations in the galaxy counts are very small and one might expect the ionizing background to be extremely uniform. However, equation (2) assumes that the source counts within each attenuation volume are Poisson distributed

around the fixed value \bar{N}_0 ; in actuality, deterministic halo clustering causes fluctuations in the *expected* number of halos within each attenuation volume, with a fractional amplitude $\sim \bar{b}\sigma(\lambda_{\text{mfp}})$ where \bar{b} is the mean bias of the star-forming galaxies. For our fiducial model at $z = 6$, we find $\bar{b} \sim 2.5$, which yields fractional variations of order ~ 2 for $\lambda_{\text{mfp}} = 30$ Mpc.

To crudely account for this deterministic clustering, we approximate the underlying density field as linear on scales $\sim \lambda_{\text{mfp}}$ (so that the PDF of density, $p(\delta)$, is a gaussian with variance $\sigma^2(\lambda_{\text{mfp}})$) and vary the expected number of sources within each attenuation volume according to linear theory, so that $\bar{N}_0 \rightarrow \bar{N}_0(\delta, \lambda_{\text{mfp}})$, which depends on the conditional mass function for a region with a fractional overdensity δ on a scale λ_{mfp} (Lacey & Cole 1993). Thus

$$\frac{dp(>X)}{dX} = \int d\delta p(\delta) \frac{dp(X|\bar{N}_0[1+\bar{b}\delta])}{dX}. \quad (4)$$

We shall see that this is only a crude approximation, for two reasons. First, our prescription for the expected source density is not terribly accurate, since the typical fractional fluctuation in the galaxy counts is not too much smaller than unity, and because we effectively assign the modified source density to all neighboring regions as well. More importantly, nonlinear clustering is actually quite important to the PDF.

2.3.2 Flux Power Spectrum

Similar techniques also allow one to compute the correlation function or power spectrum of the intensity field due to randomly distributed sources (Zuo & Phinney 1993). However, in this case it is possible to include (linear) clustering exactly by using arguments similar to the “halo” model that describes dark matter and galaxy clustering.

The intensity field around a single galaxy of mass m is

$$J(r|M_i) = \frac{\epsilon_{\text{ion}} M_i \exp(-r/\lambda_{\text{mfp}})}{(4\pi r)^2}. \quad (5)$$

In our calculations, the normalization factor is set as in the semi-numeric model. The total intensity field is simply the sum of the profiles around each galaxy in the Universe. By analogy to the halo model, the power spectrum can then be written as the sum of two terms, one arising from correlations within a single source's envelope, and one arising from correlations between the locations of two sources (Scherrer & Bertschinger 1991; Cooray & Sheth 2002). These two terms both depend on the normalized Fourier transform of the single source profile, equation (5):

$$u_J(k) = \frac{\arctan(k\lambda_{\text{mfp}})}{k\lambda_{\text{mfp}}}, \quad (6)$$

which is independent of galaxy mass (because it only describes the shape, not the amplitude, of the intensity profile). Note that $u_J \rightarrow 1$ for $k\lambda_{\text{mfp}} \ll 1$ and $u_J \rightarrow \pi/(2k\lambda_{\text{mfp}})$ for $k\lambda_{\text{mfp}} \gg 1$.

The “single-source” term, describing correlations within randomly distributed sources’ profiles, is

$$P_{1s}(k) = |u_J(k)|^2 \int_{M_{\min}}^{\infty} dm n(m) \left(\frac{m}{\bar{\rho}_{\text{gal}}} \right)^2, \quad (7)$$

where $n(m)$ is the mass function of the sources (or the halo mass function in our calculations) and $\bar{\rho}_{\text{gal}} = n_i \bar{m}_{\text{gal}}$ is the mean mass density of galaxies (and \bar{m}_{gal} is the mean mass of the ionizing sources). Note that the single-source correlation function can also be derived from the two-point probability density (Zuo & Phinney 1993; see also Furlanetto 2009). Furthermore, the amplitude is determined by the luminosity function, but the shape is entirely determined by our imposed attenuation profile.

Similarly, the term describing the correlations between two sources may be written

$$P_{2s}(k) = |u_J(k)|^2 \bar{b}^2 P_{\text{lin}}(k, z) \left[\int_{M_{\min}}^{\infty} dm n(m) \left(\frac{m}{\bar{\rho}_{\text{gal}}} \right) \right]^2, \quad (8)$$

where $P_{\text{lin}}(k, z)$ is the linear matter power spectrum. Again, this two-point function has a shape fixed by the source profiles, but an amplitude that depends on the nature of the sources. In fact, the term in square brackets is simply equal to unity; we have left it in only for comparison to the halo model, where $u(k)$ is a function of m and so must be left inside the integral. Because $|u_J|^2 \rightarrow 1$ for small k , the two-source term traces the linear matter power spectrum on these scales, before being truncated at $k\lambda_{\text{mfp}} \sim 1$.

If we take a simple model in which all sources have a fixed mass, we find that

$$\frac{P_{2s}}{P_{1s}} \approx \bar{b}^2 n_i P_{\text{lin}}(k). \quad (9)$$

Thus the two-source term becomes more and more important as the number density of sources, their bias, and the underlying linear clustering increase: this is simply because all of these increase the importance of deterministic clustering relative to the correlations inherent to the flux profile itself.

3 THE FLUX PDF

3.1 Attenuation

As discussed above, we account for flux attenuation through the factor $e^{-r/\lambda_{\text{mfp}}}/(4\pi r^2)$ in eq. (1). Thus flux attenuation is taken

into account in a spatially homogeneous way through the free parameter λ_{mfp} . Although such an approach misses the obvious correlation of absorbers (LLSs and DLAs) with sources, it had several benefits over using the typical approximate radiative transfer algorithms in cosmological hydrodynamical simulations. Firstly, our approach is versatile, easy to calibrate and to run quickly for large-scale boxes. This allows us to easily explore the effects of varying λ_{mfp} , which is important, given the large uncertainties in this value at these redshifts (e.g. Faucher-Giguère et al. 2008). Secondly, our formalism allows for easy comparisons with analytic models, and provides direct insight into the features of the UVB. Finally, and most importantly, large-scale cosmological hydrodynamical simulations *do not directly resolve LLSs and DLAs*; thus they too must employ analytic prescriptions to populate their box with absorbers. We also note that the mean free paths at these redshifts is so large ($\gtrsim 20$ Mpc; Storrie-Lombardi et al. 1994; Madau et al. 1999; Miralda-Escudé 2003; Faucher-Giguère et al. 2008) that flux received by each cell is sourced by tens of thousands of sources, likely averaging out detailed radiative transfer effects. We thus do not expect our results to be significantly affected by the fact that we neglect the details of radiative transfer and the clustering of absorbers.

Nevertheless, stochastic fluctuations in the absorber population will help to increase the variance in the radiation field: if LLSs are primarily responsible for absorption, each attenuation length will contain only a single absorber on average and the absorption length will vary quite strongly. However, in practice $\gtrsim 50\%$ of the absorption is typically due to lower-column density absorbers (see, e.g., the Appendix to Furlanetto & Oh 2005), so this is probably not a large factor for us.

In Figure 2, we plot the flux PDFs for a $L = 100$ Mpc box with 0.5 Mpc cells at $z = 10$, $\lambda_{\text{mfp}} = 10$ Mpc, and $M_{\min} = 10^8 M_{\odot}$. We show the effects of truncating the sum over contributing sources in eq. (1) at distances of $1\lambda_{\text{mfp}}$, $2\lambda_{\text{mfp}}$, and $3\lambda_{\text{mfp}}$ (*left to right solid curves*). The distributions appear to converge at $2\lambda_{\text{mfp}} \rightarrow 3\lambda_{\text{mfp}}$. The high end tails of the distribution are dominated by nearby ionizing sources and are fairly insensitive to the effective horizon (i.e. the high-value tails are roughly constant as more distant sources are included in the UVB calculation). When plotted in unnormalized units (i.e. J instead of $J/\langle J \rangle$), this overlap is more pronounced and the leftmost (blue) curve shifts to even smaller values. In contrast, the low-end tails of the PDFs correspond to locations far away from ionizing sources and so are very sensitive to the inclusion of more distant sources, which roughly translates to changing the amplitude an effective homogeneous UVB. If sources were homogeneously distributed, truncating the flux distribution at distances greater than $n\lambda_{\text{mfp}}$ would underestimate the ionizing flux by a factor of $(1 - e^{-n})$. The clustering of sources results in the more complicated behavior seen in Fig. 2.

The dotted curve in Fig. 2 neglects the exponential attenuation term in eq. (1), and instead only includes the contributions of sources at distances $r < \lambda_{\text{mfp}}$. This is analogous to assuming a step-function optical depth so that $\tau = 0$ for distances $r < \lambda_{\text{mfp}}$ and $\tau = \infty$ at distances $r > \lambda_{\text{mfp}}$. Such a prescription can simplify computation in analytic models (e.g., Zuo 1992) and numerical calculations (e.g., Bolton et al. 2006) and seems to result in fairly accurate flux distributions, though it slightly overestimates the number of faint regions. This overestimate worsens somewhat if one integrates sources only out to $1\lambda_{\text{mfp}}$ but also includes exponential attenuation (e.g. Santos et al. 2008).

In Figure 3, we explore the impact of varying λ_{mfp} . Flux PDFs at $z = 5$ (*solid curves*) and $z = 6$ (*dashed curves*) are shown as

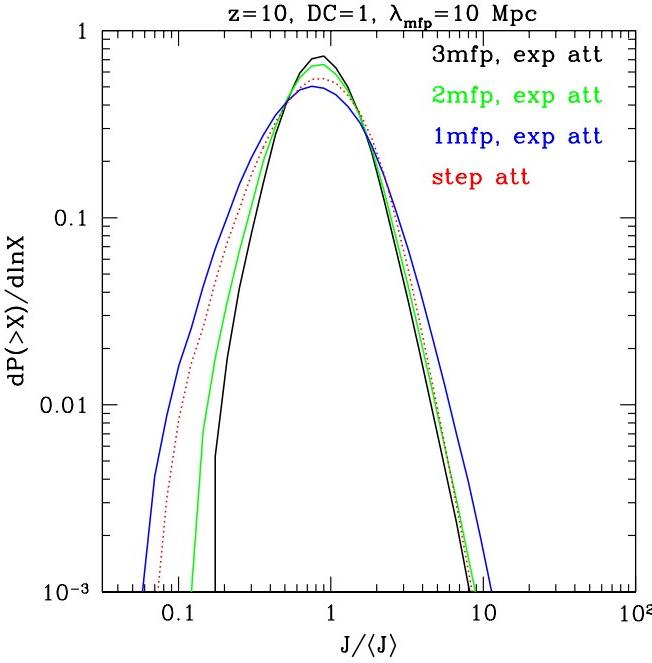


Figure 2. Flux PDFs at $z = 10$, assuming $\lambda_{\text{mfp}} = 10 \text{ Mpc}$, $M_{\text{min}} = 10^8 M_\odot$ and DC=1.0, calculated according to eq. (1) but truncating the sum over contributing sources at distances of $1\lambda_{\text{mfp}}$, $2\lambda_{\text{mfp}}$, and $3\lambda_{\text{mfp}}$ (left to right solid curves). The dotted curve assumes a step-function attenuation, essentially $\tau = 0$ for distances $r < \lambda_{\text{mfp}}$ and $\tau = \infty$ at distances $r > \lambda_{\text{mfp}}$.

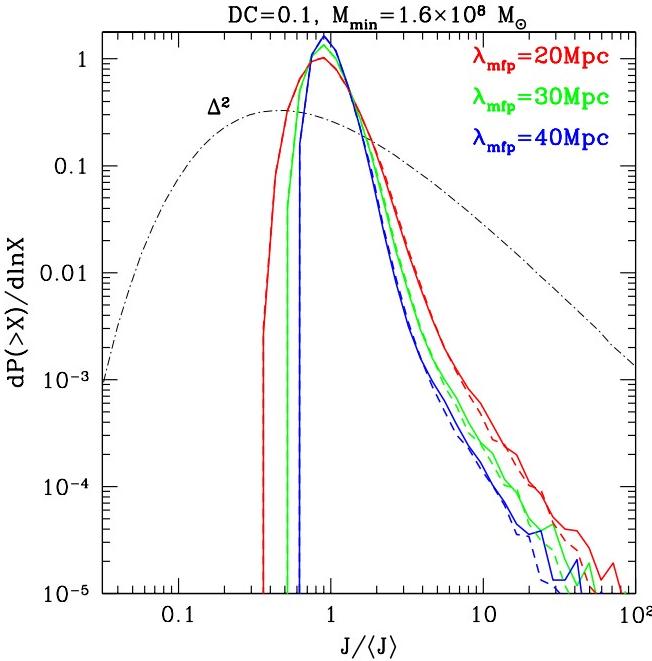


Figure 3. Flux PDFs at $z = 5$ (solid curves) and $z = 6$ (dashed curves) generated assuming $M_{\text{min}} = 1.6 \times 10^8 M_\odot$ and DC=0.1. Starting from the left side of the figure, curves correspond to values of $\lambda_{\text{mfp}} = 20, 30, 40 \text{ Mpc}$, respectively. For comparison, the dot-dashed curve shows the squared overdensity, $\Delta^2 \equiv (\rho/\bar{\rho})^2$, taken from the hydrodynamical simulation at $z = 5.71$ with cell size 0.74 Mpc.

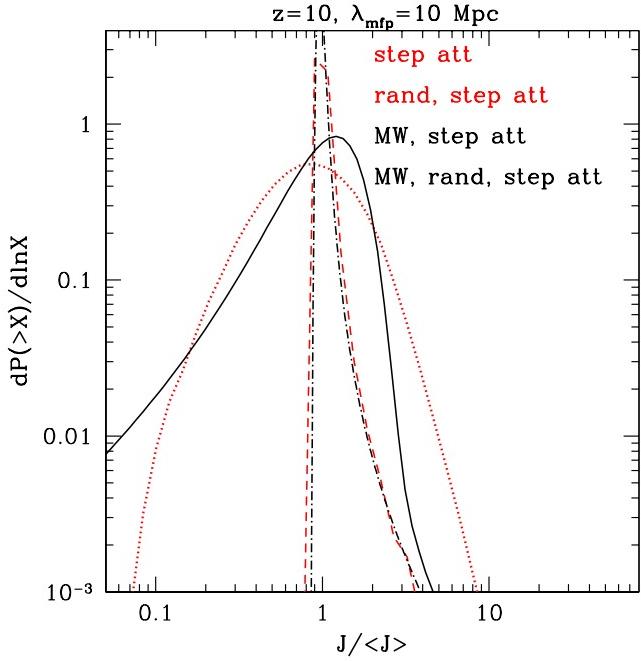


Figure 4. Flux PDFs at $z = 10$, assuming $\lambda_{\text{mfp}} = 10 \text{ Mpc}$, $M_{\text{min}} = 10^8 M_\odot$ and DC=1.0. The dotted curve was generated assuming a step-function attenuation, essentially $\tau = 0$ for distances $r < \lambda_{\text{mfp}}$ and $\tau = \infty$ at distances $r > \lambda_{\text{mfp}}$. The dashed curve is obtained in the same manner, but with randomized source locations (i.e. ignoring clustering). The solid and dot-dashed curves correspond to the analytical model described in the text, including and ignoring clustering on scales $> \lambda_{\text{mfp}}$ as described in equation (4), respectively.

suming $M_{\text{min}} = 1.6 \times 10^8 M_\odot$ and DC=0.1. Note that the evolution of structure from $z = 10$ in Fig. 2 to $z = 5$ and 6 in Fig. 3 preferentially suppresses the low-end tail of the PDFs as regions containing few ionizing sources become extremely rare. This results in narrower and more asymmetric distributions.

As expected, increasing the m.f.p. also results in narrower flux distributions, as the gas cells see increasingly representative volumes of space inside their horizon. We note however that the $\lambda_{\text{mfp}} = 40 \text{ Mpc}$ distributions are still fairly broad. Understandably, when plotted in unnormalized units (i.e. J instead of $J/\langle J \rangle$), the high-value tails overlap as they are dominated by nearby halos. Perhaps surprisingly, we see very little variation of the flux PDF with redshift, over the range $z = 5$ to 6.

3.2 Source Clustering

The clustering of sources broadens the flux PDF and creates a high-value tail in the distribution (e.g. Mesinger & Dijkstra 2008; Dijkstra et al. 2008). How strong is this effect and how much do the properties of the underlying sources affect these trends?

In Fig. 4, we plot flux PDFs at $z = 10$, assuming $\lambda_{\text{mfp}} = 10 \text{ Mpc}$. The dotted curve in the figure was generated from the semi-numeric calculation assuming a step-function attenuation at $\lambda_{\text{mfp}} = 10 \text{ Mpc}$, shown to be an adequate approximation in Fig. 2. The dashed curve is obtained in the same manner, but with randomized source locations. We see that ignoring source clustering can severely underestimate the widths of the flux PDF. Even the shape of the curves is quite different.

Fig. 4 also compares the PDF to our analytic model. The dot-

dashed line uses the analog of equation (2) but for step-wise attenuation (see Zuo 1992 and Furlanetto 2008), again assuming Poisson distributed sources. It closely matches the semi-numeric results for the same assumptions.

The solid line includes linear clustering over the scale λ_{mfp} as described in equation (4). The clustering length of $M \sim M_{\min} \sim 10^8 M_\odot$ halos is ~ 2 Mpc at this redshift, so the clustering on scales larger than the m.f.p. should be well into the linear regime and accurately predicted by the analytic model. Obviously, this simple prescription provides a relatively poor match to the full semi-numeric results: although the large-scale clustering does broaden the distribution by a comparable amount, the shapes still disagree.

Part of this difference is easy to understand and is a consequence of the analytic model’s simple prescription for clustering: the model not only fixes the *local* attenuation volume to have a given density, but it also effectively assumes that surrounding regions have the same number of sources. This artificially suppresses the contribution from distant galaxies (which provide 2/3 of the mean ionizing background in these models). In practice, points inside low density regions should be exposed to higher flux values from neighboring higher density regions. Such long-wavelength effects are missed in the simple analytic treatment, resulting in the spuriously long small-value tail in the PDF.

However, by the same logic our prescription places high-density regions near high-density neighbors, so one might expect us to overestimate the extent of the variation at high J as well – but in fact the prescription *underestimates* the width there as well. This is because, in contrast to the low- J points, the radiation field at such high levels is not dominated by the background of distant sources but by nearby clusters of objects, for which smaller-scale clustering and nonlinear effects are important.

This is not particularly surprising in light of the analogy we have drawn to the halo model: although that simple prescription does a remarkably good job of matching the dark matter power spectrum, it cannot easily be used to predict the actual density distribution of the IGM. Indeed, Scherrer & Bertschinger (1991) have shown that the PDF of the fractional deviation of J from the mean value, $1 + \delta_J = J/\langle J \rangle$, is

$$\frac{dp(>\delta_J)}{d\delta_J} = \int_{-\infty}^{\infty} \frac{ds}{2\pi} e^{-i\delta_J s} \exp \left[\sum_{p=1}^{\infty} \frac{(is)^p}{p!} \xi_J^{(p)}(0, 0, \dots, 0) \right], \quad (10)$$

where $\xi_J^{(p)}(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_p)$ is the p -point correlation function of the radiation field. Thus the PDF depends on *all* higher order correlations at zero-lag, which are evidently strongly affected by the nonlinear clustering on small scales. The simple addition of large-scale linear clustering is not sufficient to match the semi-numeric results.

To see the impact of the source properties on the clustering signature in the flux PDF, in Fig. 5, we plot several flux PDFs assuming $\lambda_{\text{mfp}} = 30$ Mpc. The dotted, short-dashed and solid curves were all generated at $z = 5$, but assuming various values of DC and M_{\min} . These two parameters drive the correlation function of sources and are crucial in constraining the UV luminosity to halo mass relation at high- z (e.g. Lee et al. 2008). From Fig. 5, we see that the low-end cut-off is generally unaffected by these parameters. This is to be expected since, as we mention above, the low-end tail of the flux distribution corresponds to voids distant from UV sources, and are sensitive to the effective horizon λ_{mfp} (which is identical in all these curves). Thus, most of the clustering imprint is in the high-value tails of the PDFs.

The figure also confirms that the widest flux PDF results from the model in which sources are the most biased and rare: having

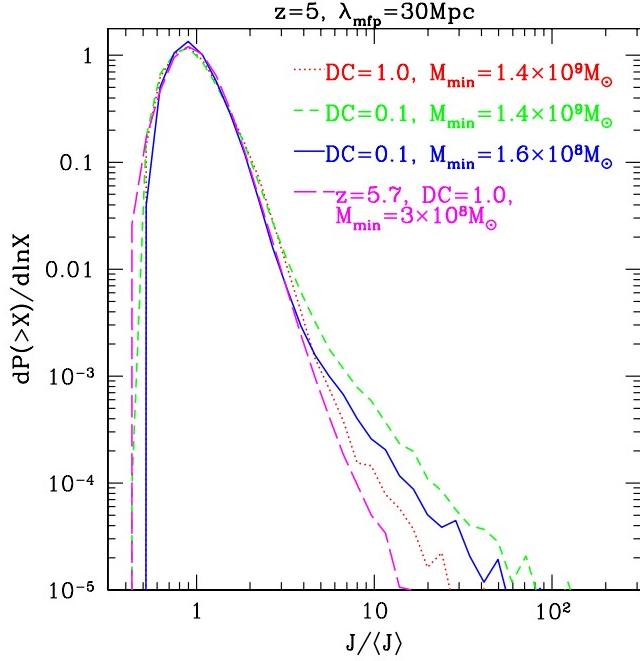


Figure 5. Flux PDFs assuming $\lambda_{\text{mfp}} = 30$ Mpc. Dotted, short-dashed and solid curves correspond to $z = 5$ and $(\text{DC}, M_{\min}/M_\odot) = (1.0, 1.4 \times 10^9)$, $(0.1, 1.4 \times 10^9)$, and $(0.1, 1.6 \times 10^8)$, respectively. The long-dashed curve is generated from the $z = 5.71$ star-formation rate output from the numerical results of Trac et al. (2008), as described in the text.

the highest M_{\min} and lowest DC. As the duty cycle is increased, the high-value tail retracts. A similar trend is seen as one lowers M_{\min} .

As mentioned above, the choice of $M_{\min} = 1.4 \times 10^9 M_\odot$ was motivated by the fact that $(\text{DC}, M_{\min}/M_\odot) = (1.0, 1.4 \times 10^9)$ and $(0.1, 1.6 \times 10^8)$ result in the same number density of sources, and thus are useful in isolating the imprint of each these two parameters. Even in our simple model where the duty cycle is not a function of mass, these corresponding curves are quite similar. Thus it is unlikely that the flux PDF, even if it could be measured, could provide a useful constraint on these parameters.

The long-dashed curve in Fig. 5 shows the flux PDF obtained from the $z = 5.71$ star-formation rate output from the hydrodynamical simulation of Trac et al. (2008), as described in §2.2. Despite differences in the star-formation prescription, this curve agrees well with the above trends. Note that the slight offset at low J is attributable to the fact that these fluxes were calculated including sources within $1.5\lambda_{\text{mfp}}$, instead of within $2\lambda_{\text{mfp}}$, as was the case for the semi-numerical boxes.

4 THE FLUX POWER SPECTRUM

4.1 Analytic Results

We begin our discussion of the flux power spectrum by examining the “halo model” analytic results (which we expect to be more successful than the PDF). Figure 6 shows some examples at $z = 10$; here $\Delta^2(k, z) = k^3/(2\pi^2)P(k, z)$. The solid curve takes $M_{\min} = 10^8 M_\odot$, DC = 1, and $\lambda_{\text{mfp}} = 10$ Mpc; the dotted curves show the corresponding one- and two-source contributions. Clearly these two terms are cleanly separated in scale for this set of parameters. The two-source term dominates on large scales;

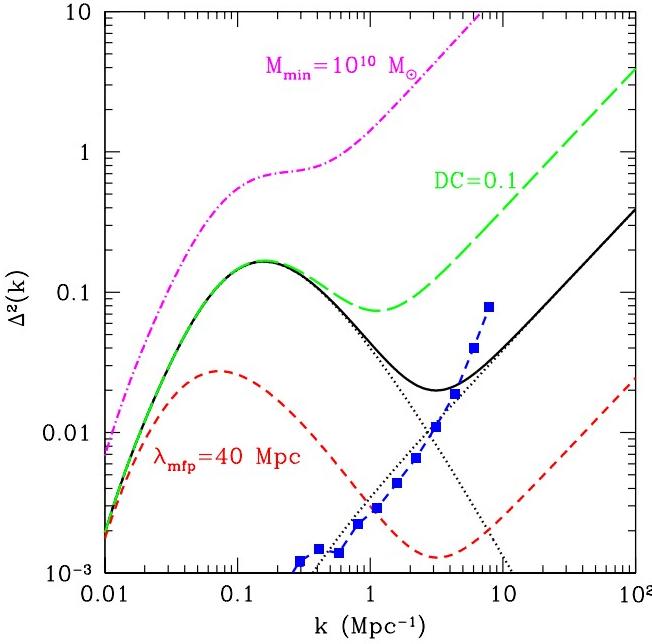


Figure 6. Flux power spectra at $z = 10$ computed with our analytic model. The solid black line takes $M_{\min} = 10^8 M_\odot$, DC= 1, and $\lambda_{\text{mfp}} = 10$ Mpc; the dotted curves show the corresponding one- and two-source contributions. The curve with solid squares shows the corresponding one-source power spectrum computed from the semi-numeric simulations. The long-dashed line takes DC= 0.1 instead, and the short-dashed line takes $\lambda_{\text{mfp}} = 40$ Mpc. The dot-dashed line takes $M_{\min} = 10^{10} M_\odot$.

for $k\lambda_{\text{mfp}} \ll 1$, it traces the linear power spectrum with an amplitude determined by the mean galaxy bias. It is then damped at $k\lambda_{\text{mfp}} \sim 1$: the short-dashed curve, which takes $\lambda_{\text{mfp}} = 40$ Mpc shows how the turnover scale shifts with the attenuation length.

On small scales, the one-source term dominates: it has $\Delta_{1s}^2 \propto k^3 u_J(k)^2 \propto k/\lambda_{\text{mfp}}^2$. Thus its amplitude decreases rapidly with the attenuation length – although this dependence is hard to separate from the dependence on overall source density. The dashed curve with solid squares shows the one-source term computed from the semi-numeric simulations (with sources randomly distributed, so that the two-source term disappears). It provides an excellent match to the analytic model, except on scales near $k \sim 10$ Mpc $^{-1}$, where the discrete grid cells affect the semi-numeric power spectrum.

The dot-dashed curve takes a larger M_{\min} . Note that the power spectrum increases with the halo bias on very large scales. Also, the decreased source number density strongly increases the importance of the one-source term, giving a much less clean separation between the one-source and two-source regime.

4.2 Numerical Results

We now compare these analytic predictions to the semi-numerical simulations. For the numerical calculations, we use the convention $\Delta^2(k, z) = k^3/(2\pi^2 V) \langle |\delta_J(\mathbf{k}, z)|^2 \rangle_k$. In Fig. 7, we plot the flux power spectra at $z = 6$ for both the analytic (*solid curves*) and semi-numeric (*dashed curves*) models, generated assuming $M_{\min} = 1.6 \times 10^8 M_\odot$ and DC=0.1. Curves correspond to values of $\lambda_{\text{mfp}} = 20, 30, 40$ Mpc, from top to bottom.

The agreement between the two calculations is quite impres-

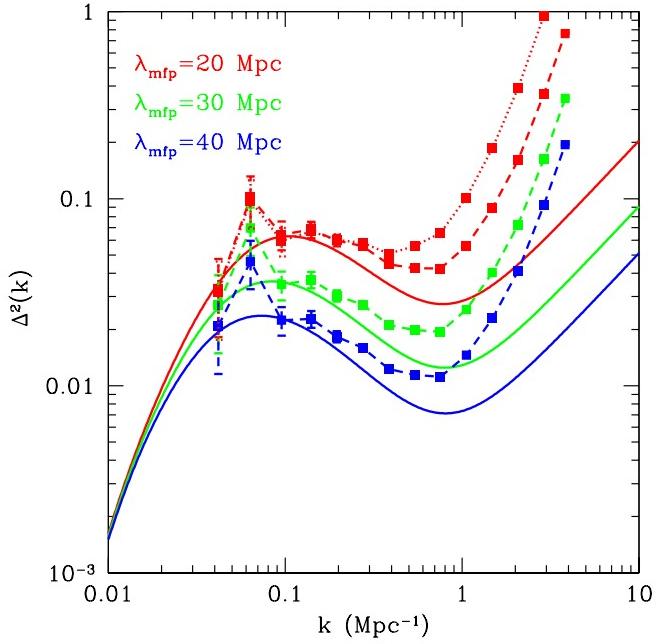


Figure 7. Flux power spectra at $z = 6$ from the analytic model (*solid curves*) and the semi-numeric simulations (*dashed curves*, with Poissonian error bars). All curves assume $M_{\min} = 1.6 \times 10^8 M_\odot$ and DC=0.; the three sets correspond to values of $\lambda_{\text{mfp}} = 20, 30, 40$ Mpc, from top to bottom. The dotted curve shows the semi-numeric result at $z = 5$ for $\lambda_{\text{mfp}} = 20$ Mpc.

sive, especially at large scales. The analytic model clearly shows the separation between the two-source and single-source terms; with the semi-numeric results the difference is less clear because the boxes are not large enough to sample the regime $k\lambda_{\text{mfp}} \gg 1$, where Δ^2 traces the matter power spectrum. However, the amplitudes of the analytic and semi-numeric calculations match reasonably well on these scales: the differences in the two-source regime are attributable to slightly different mean halo biases in the analytic and semi-numeric calculations.

However, the semi-numeric results have a steeper slope for the one-source component; this is most likely due to nonlinear clustering, which will steepen the correlation function on small scales, since Fig. 6 shows that the two different methods agree very well for the one-source term when clustering is ignored. Interestingly, however, the crossover scale between the one-source and two-source regimes is accurately predicted by the analytic model.

The three sets of curves in Fig. 7 show that the m.f.p. has little impact in the shape of the spectrum over the scales probed by the semi-numeric simulations: it primarily serves to decrease power fairly evenly on all scales with $k\lambda_{\text{mfp}} \lesssim 1$. As described above, for smaller wave-numbers the attenuation length has no effect: effectively, a larger attenuation length causes more damping relative to the linear power spectrum.

Finally, the dashed curve shows the flux power spectrum from the semi-numeric simulations at $z = 5$ and $\lambda_{\text{mfp}} = 20$ Mpc. By comparing Fig. 7 and Fig. 3, we see that the flux power spectra are more sensitive to redshift than the flux PDFs. The differences are similar for other values of λ_{mfp} , and in the analytic model: in all cases, the amplitude of the one-source component is somewhat larger at the lower redshift.

This is easily understood from the analytic formalism. The

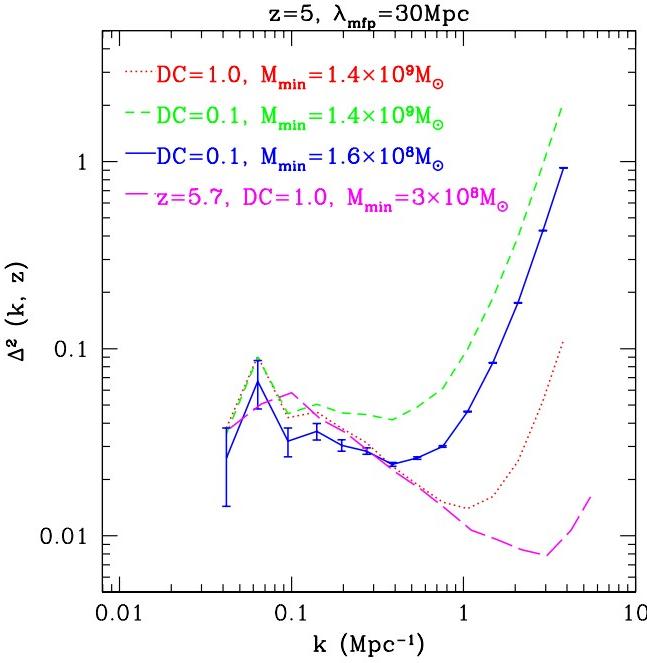


Figure 8. Flux power spectra analogous to the models shown in Fig. 5.

two source term is proportional to $\bar{b}^2 P_{\text{lin}}(k, z)$; each of these terms evolves with redshift, but in the product the evolution nearly cancels at high redshift (Oh 2001). The one source term is proportional to the second moment of the mass function (when integrated over mass – eq. 7), which increases toward lower redshifts as more and more massive halos are assembled.

A similar argument holds for the power spectra presented in Fig. 8, where we vary the DC and M_{\min} parameters. The curves correspond to the same models presented in Fig. 5. Here we see drastic differences on moderate to small scales, because models in which the sources are rare (i.e. with low DC and high M_{\min}) transition to the single-source regime on larger scales. If observed, this statistic could be a powerful tool in constraining these parameters. The flux field has the benefit of encoding information about ionizing sources which are too faint to be detected by galaxy surveys. Unfortunately, such flux power spectra will be nearly impossible to extract from the Ly α forest, even given a very large number of sightlines (see Fig. 11 and associated discussion).

5 LYMAN α FOREST

The most obvious application of our flux field analysis is in the interpretation of the Ly α forest spectra following reionization. Thus far, almost all estimates of the photoionization rate per hydrogen atom, $\Gamma_{12} \equiv \Gamma / 10^{-12} \text{ s}^{-1}$, obtained from the Ly α forest have assumed a uniform UV background (e.g. Fan et al. 2006). However, in §3 we showed that UV background actually has sizable fluctuations in this regime. These fluctuations could potentially introduce bias and stochasticity in the determinations of Γ_{12} from Ly α forest attenuation measurements.

Bolton & Haehnelt (2007) briefly addressed this issue by including a similar UV background model to our semi-numeric approach. They confirm the results of lower redshift studies (Gnedin & Hamilton 2002; Meiksin & White 2004; Bolton et al. 2006) that assuming a homogeneous UVB underestimates the mean

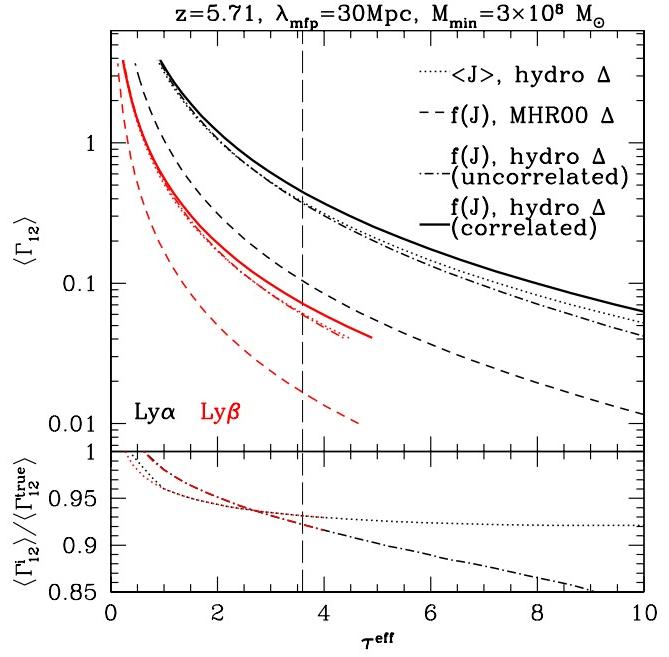


Figure 9. *Top panel:* Mean ionization rate corresponding to an effective Ly α (black curves) and Ly β (red curves) GP optical depth, for the $z = 5.71$ hydrodynamical simulation box of Trac et al. (2008), assuming $\lambda_{\text{mfp}} = 30$ Mpc. Dotted curves use the density field from the simulation, but assume a uniform UVB. Dashed curves assume a fluctuating UVB from our simulations, but convolved with the Miralda-Escudé et al. (2000) density distribution. Dot-dashed curves use the fluctuating UVB and density field from the simulations, but randomly correlate the two. Finally, solid curves use the correlated fluctuating UVB field and underlying density field from the simulations. *Bottom panel:* Ratio of the mean ionization rate for each model i and the “true” model shown as the solid curves in the top panel. The vertical dashed line demarcates the average value of $\tau_{\alpha}^{\text{eff}} \sim 3.6$, roughly corresponding to the value measured from the SDSS QSOs at this redshift (taken from Fan et al. 2006).

ionization rate obtained from Ly α forest absorption, $\langle \Gamma_{12} \rangle$. They find that this effect is of order $\sim 1\% - 10\%$, depending on λ_{mfp} and M_{\min} (see their Fig. 3). However, for this analysis they used a single realization of a $30 \text{ h}^{-1} \text{ Mpc}$ simulation box. This box size is comparable to a single ionizing photon m.f.p at these redshifts (Storrie-Lombardi et al. 1994; Madau et al. 1999; Miralda-Escudé 2003; Faucher-Giguère et al. 2008) and is several times smaller than the length of the GP troughs used to estimate $\langle \Gamma_{12} \rangle$. Thus, their analysis was unable to fully take into account cosmic variance.

Using our large simulation boxes, we wish to systematically investigate this effect. As a first step, we do not create detailed mock absorption spectra. Instead, for a given flux field, we merely calculate the Ly α optical depth, τ_{α} , for each cell, assuming ionization equilibrium. Then we average the flux decrements over the entire box to obtain a corresponding effective GP optical depth: $\tau_{\alpha}^{\text{eff}} \equiv -\ln \langle \exp[-\tau_{\alpha}] \rangle_x$.

In the top panel of Fig. 9 we plot $\langle \Gamma_{12} \rangle$ for each corresponding effective Ly α (black curves) and Ly β (red curves) GP optical depth, for the $z = 5.71$ hydrodynamical simulation box, assuming $\lambda_{\text{mfp}} = 30$ Mpc. The cell size for this simulation is 0.74 Mpc . As mentioned above, the density field was calculated on a much finer grid of 0.19 Mpc , which resolves the Jeans length in the mean density, ionized IGM by a factor of few. The smoothed cell-size of

0.74 Mpc corresponds to the Jeans mass in the mean density ionized IGM, and is roughly comparable to the Doppler width of the Ly α line at $T = 10^4$ K, as well as to the Keck ESI resolution at this redshift. Thus we expect our $\tau^{\text{eff}} \leftrightarrow \langle \Gamma_{12} \rangle$ mapping to be adequate [indeed, our results are consistent with the mock spectra analysis of Bolton & Haehnelt (2007)], but we stress that we did not create detailed mock spectra.

The curves in Fig. 9 show the effects of various simplifications on the value of $\langle \Gamma_{12} \rangle$ inferred from the τ^{eff} measured in the spectra. The most accurate model, using both our fluctuating flux fields and underlying density fields, is shown by the solid curves. By comparing the solid and dotted curves (or their ratios in the bottom panel), we see that assuming a uniform UVB makes very little impact on the estimated mean ionization rate per hydrogen atom. At low values of τ^{eff} , the two curves are practically indistinguishable, with the difference increasing with τ^{eff} . At $\tau_\alpha \sim 3.6$ (roughly corresponding to the observed value at this redshift, e.g. Fan et al. 2006), assuming a uniform UVB results in a $\sim 5\%$ underestimate in the value of $\langle \Gamma_{12} \rangle$.

It is fairly easy to interpret this behavior by looking back at Fig. 3. The optical depth approximately scales as $\tau \propto \Delta^2/\Gamma$, where $\Delta \equiv \rho/\bar{\rho}$ is the gas density in units of the mean. Thus the optical depth distribution can be approximately obtained by convolving these two quantities. From Fig. 3, we see that the distribution of Δ^2 is much wider than all of the flux distributions. Thus a uniform UVB, which treats the flux distributions as a delta function, is an accurate assumption. Because of the exponential attenuation involved, the τ^{eff} statistic is sensitive to the low-end tail of the density distribution. As the ionizing background gets fainter (i.e. increasing τ^{eff}), the τ^{eff} statistic is sensitive to an increasingly narrow range in the low-end tail of the optical depth distribution. It is in this regime that the non-zero width of the ionizing background PDF becomes evident, as the convolution with Δ^2 is mostly evident in smearing-out this low-end tail. Thus assuming a uniform UVB underestimates the true $\langle \Gamma_{12} \rangle$, but only slightly, with this underestimate increasing with increasing τ^{eff} .

Figure 9 also shows the impact of other common assumptions in the $\tau^{\text{eff}} \leftrightarrow \langle \Gamma_{12} \rangle$ mapping. The dot-dashed curve is generated by using the same UVB and density fields as in the solid curve, but randomly pairing the two. This also results in a slight underestimate of $\langle \Gamma_{12} \rangle$, though a little more notable than the uniform UVB assumption. Since underdense pixels are preferentially farther away from most sources, where the UVB is lower (c.f. Fig. 1), not correlating the two fields results in underdense pixels receiving fluxes which are too strong, thus lowering $\langle \Gamma_{12} \rangle$. Nevertheless, the effect is still small because of the relative width of the PDFs argument in the above paragraph.

Finally, the most drastic effect, resulting in an underestimate of a factor of ~ 2 in our models, comes from using the analytic density PDF from Miralda-Escudé et al. (2000). This offset is normally attributed to errors in the fluctuating Gunn-Peterson approximation, which ignores effects such as the line wings, peculiar velocities, and blending. However, our numeric results make a similar approximation, so that cannot be the source of the error here. Instead, it may be differences between the assumed density distribution (generated to fit numerical simulations at redshifts $z \sim 2-4$, using a simulation box of only $10 h^{-1}$ Mpc) and that of the simulation. At higher redshifts (i.e. τ^{eff} values), the transmission becomes dominated by rare voids (Oh & Furlanetto 2005; Lidz et al. 2006), and so cosmic variance is very important. Proper analysis requires large simulation boxes with accurate hydrodynamical treatment and direct higher-redshift outputs. However, if the semi-analytic method

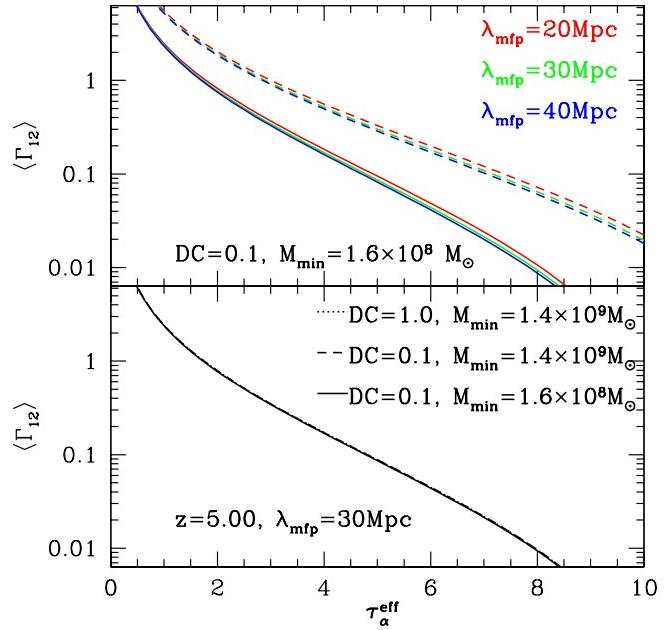


Figure 10. $\tau^{\text{eff}} \leftrightarrow \langle \Gamma_{12} \rangle$ mappings using our semi-numerical simulations, assuming various flux models. In the top panel, solid curves correspond to $z = 5$, while dashed curves correspond to $z = 6$.

is renormalized using a simulation (as is typical), the *shapes* of the curves are similar enough that the errors will be relatively small, so the approximation is still useful for understanding the evolution of $\langle \Gamma_{12} \rangle$. Furthermore, it is not clear how this offset is affected by the assumed thermal history of the gas, which sets the effective local smoothing scale. A detailed exploration of this is beyond the scope of this work, and merits its own study.

In Fig. 10, we show the $\tau^{\text{eff}} \leftrightarrow \langle \Gamma_{12} \rangle$ mappings using our semi-numeric simulations, assuming various flux models. Note that the semi-numerical density fields slightly overpredict the presence of very rare voids which dominate this statistic (Mesinger et al., in preparation), thereby over-predicting the transmission for a given UVB. It is remarkable how insensitive this mapping is to the details of the UVB, despite the notable differences in the flux PDFs and power spectra shown in the previous section. Again, this is due to the fact that the fluctuations in the UVB, and the differences among various models of the UVB, are overshadowed by the fluctuations in the density field. The only notable differences among the curves is due to the density field evolution from $z = 6$ to $z = 5$.

We note that our results are in general agreement with Bolton & Haehnelt (2007), who found that assuming a homogeneous UVB underestimates $\langle \Gamma_{12} \rangle$, by a few percent over this range of λ_{mfp} values. However, they find a somewhat stronger dependence on λ_{mfp} than we do. This is most likely a result of cosmic variance attributable to the small simulation box they used to study this effect.

Finally, we wish to quantify how easy it is to distinguish various UVB models from observations of the Ly α forest. One can see from Fig. 8 that the shape of the flux power spectrum encodes information about the clustering properties of the sources. As mentioned above, the flux field (as potentially extracted from the Ly α forest), has a great benefit in being able to probe the impact of ionizing sources below the detection thresholds of galaxy surveys. Hence, in Fig. 11 we show the optical depth power spectra at $z = 5$ and

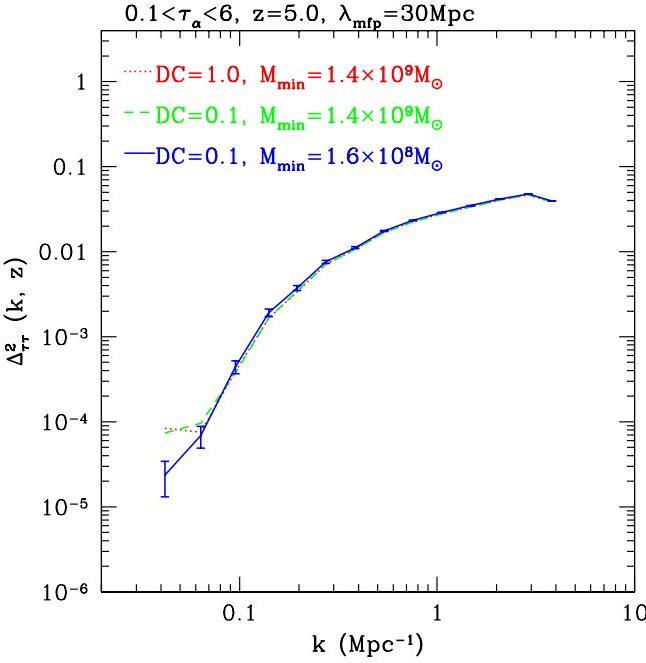


Figure 11. Optical depth power spectra at $z = 5$ and $\lambda_{\text{mfp}} = 30$ Mpc, within the range $0.1 < \tau_\alpha < 6$, approximately corresponding to the range probed by current instruments.

$\lambda_{\text{mfp}} = 30$ Mpc for some of the same models as in Fig. 8. Unfortunately, the optical depth power spectra of these models are practically indistinguishable from each other, despite the large differences shown in Fig. 8. This result could have been foreshadowed from Fig. 10 and is again due to the fact that the density field and its fluctuations swamp the UVB signal.

However, we note that the density fluctuations decrease as redshift increases, while fluctuations in the UVB increase. The UVB becomes more inhomogeneous not only through the evolution of λ_{mfp} and the bias of sources, but it is also modulated by the ionized bubbles throughout reionization (whenever that epoch may be). Thus inhomogeneities in the flux field become increasingly important at higher redshifts, although whether this can be observed in the highly saturated forest at $z \gtrsim 6$, where the density fluctuations themselves cause substantial fluctuations in τ^{eff} (Lidz et al. 2006), is questionable.

6 CONCLUSIONS

We systematically study the expected fluctuations in the hydrogen ionizing background in the epoch following reionization ($z \sim 5$ –6). Unlike at lower redshifts ($z \lesssim 4$), the UVB can be quite inhomogeneous in this regime. This is due to the smaller mean free path of ionizing photons, as well as the clustering of increasingly rare ionizing sources.

We confirm that there is a sizable spread (the widths of the PDFs at half of the maximum likelihoods spanning factors of ~ 2 –4) and asymmetry in the distribution of ionizing fluxes in a cosmological volume. Expected values of the m.f.p. are large enough ($\lambda_{\text{mfp}} \gtrsim 20$ Mpc), and sources are ubiquitous enough, that the PDF of the ionizing background is asymmetric, with a tail extending to higher values. This high-value tail is set by clustering on small scales, and is insensitive to the contributions of distant sources (i.e.

λ_{mfp}) (Mesinger & Dijkstra 2008). We also find that the PDFs are most sensitive to the m.f.p., with higher values of λ_{mfp} increasingly truncating the low-value tails. Furthermore, we note that the analytic formalism of Meiksin & White (2004) severely underestimates the width of the flux PDF at this epoch, since it ignores the clustering of sources.

The power spectra of the ionizing background are even more sensitive to the details of the source model. We have shown that an approach analogous to the halo model, in which contributions from the proximity zones of single sources (which dominate fluctuations on small scales and are independent of the large-scale galaxy distribution) and multiple sources (which dominate on large scales, and which depend on the large-scale clustering of the ionizing sources) are separated, provides a good match to the simulated results, although nonlinear clustering provides even more small-scale power than the analytic model predicts. The single-source component is particularly sensitive to the number density of galaxies; unfortunately, these power spectra will be nearly impossible to observe with the Ly α forest.

We also model the imprint of various UVBs in the Ly α forest at $z \sim 5$ –6. We find that the Ly α forest spectra are extremely insensitive to the details of the UVB, despite marked differences in the PDFs and power spectra of our various UVBs. This is attributable to the fact that the Ly α optical depth scales as $\tau \propto \Delta^2/\Gamma$, and the squared overdensity distribution, $f(\Delta^2)$, is much wider than all of our Γ distributions. In fact, even the extreme scenario of a uniform background (i.e. a delta function Γ distribution), only underestimates the value of $\langle \Gamma_{12} \rangle$ inferred from the Ly α forest by a few percent. Thus our results justify the common assumption of a uniform UVB in Ly α forest analysis.

We also confirm that accurate modeling of the density field is very important in the $\tau^{\text{eff}} \leftrightarrow \langle \Gamma_{12} \rangle$ relation. We find significant differences (factor of ~ 2) in this relation among the numerical and extrapolated Miralda-Escudé et al. (2000) density distributions (though we stress again that we did not create detailed mock spectra in this work). At higher redshifts, the τ^{eff} statistic increasingly depends on a narrower range in the low-value density PDF. Large-scale hydrodynamical modeling of the statistics of these rare voids is therefore essential in interpreting the high-redshift Ly α forest. Pawlik et al. (2009) have recently studied the IGM density field during reionization, and do find significant differences from the Miralda-Escudé et al. (2000) fitting formula, although their simulation boxes are much too small to study the rare voids pertinent to these efforts.

An interesting comparison can be made between our analysis of hydrogen reionization and similar effects during and after helium reionization (Dixon & Furlanetto, in preparation). In the latter case, the PDF of J is much broader (spanning over an order of magnitude), and the fluctuations have a much larger effect on the HeII Ly α forest, introducing a factor ~ 2 bias in estimates of $\langle \Gamma \rangle$ when the background is assumed uniform. This has important implications for the HeII Ly α forest during reionization: unlike the case we have studied here, the large fluctuations in the ionizing background lead to substantial opacity gaps even during the middle phases of helium reionization.

We thank Hy Trac and Renyue Cen for the use of the $z = 5.71$ density and source output of their simulations, as well as for stimulating conversations. We also thank James Bolton for helpful com-

ments on the manuscript. Support for this work was provided by NASA through Hubble Fellowship grant #HF-01222.01 to AM, awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS 5-26555. SF was partially supported by the NSF through grant AST-0607470 and by the David and Lucile Packard Foundation.

REFERENCES

- Bolton J. S., Haehnelt M. G., 2007, MNRAS, 382, 325
 Bolton J. S., Haehnelt M. G., Viel M., Carswell R. F., 2006, MNRAS, 366, 1378
 Bouwens R. J., Illingworth G. D., Blakeslee J. P., Franx M., 2006, ApJ, 653, 53
 Cooray A., Sheth R., 2002, Physics Reports, 372, 1
 Croft R. A. C., 2004, ApJ, 610, 642
 Dijkstra M., Haiman Z., Mesinger A., Wyithe J. S. B., 2008, MNRAS, 391, 1961
 Dijkstra M., Wyithe J. S. B., Haiman Z., 2007, MNRAS, 379, 253
 Dunkley J., et al., 2008, ArXiv e-prints, 803
 Fan X., et al., 2006, AJ, 132, 117
 Fardal M. A., Shull J. M., 1993, ApJ, 415, 524
 Faucher-Giguère C.-A., Lidz A., Hernquist L., Zaldarriaga M., 2008, ApJ, 688, 85
 Feng L., Bi H., Liu J., Fang L.-Z., 2008, MNRAS, 383, 1459
 Furlanetto S., 2008, ApJ, submitted (arXiv.org/0812.3411 [astro-ph])
 Furlanetto S., 2009, ApJ, submitted (arXiv.org/0901.2584 [astro-ph])
 Furlanetto S. R., Oh S. P., 2005, MNRAS, 363, 1031
 Gnedin N. Y., Hamilton A. J. S., 2002, MNRAS, 334, 107
 Kashikawa N., et al., 2006, ApJ, 648, 7
 Komatsu E., et al., 2008, ArXiv e-prints, 803
 Lacey C., Cole S., 1993, MNRAS, 262, 627
 Lee K., Giavalisco M., Conroy C., Wechsler R. H., Ferguson H. C., Somerville R. S., Dickinson M. E., Urry C. M., 2008, ArXiv e-prints
 Lidz A., Oh S. P., Furlanetto S. R., 2006, ApJ, 639, L47
 Madau P., Haardt F., Rees M. J., 1999, ApJ, 514, 648
 McQuinn M., Hernquist L., Zaldarriaga M., Dutta S., 2007, MNRAS, 381, 75
 Meiksin A., White M., 2004, MNRAS, 350, 1107
 Meiksin A. A., 2007, ArXiv e-prints
 Mesinger A., Dijkstra M., 2008, MNRAS, 390, 1071
 Mesinger A., Furlanetto S., 2007, ApJ, 669, 663
 Miralda-Escudé J., 2003, ApJ, 597, 66
 Miralda-Escudé J., Haehnelt M., Rees M. J., 2000, ApJ, 530, 1
 Oh S. P., 2001, ApJ, 553, 499
 Oh S. P., Furlanetto S. R., 2005, ApJ, 620, L9
 Okamoto T., Gao L., Theuns T., 2008, MNRAS, 390, 920
 Pawlik A. H., Schaye J., 2008, ArXiv e-prints
 Pawlik A. H., Schaye J., van Scherpenzeel E., 2009, MNRAS, 394, 1812
 Santos M. G., Amblard A., Pritchard J., Trac H., Cen R., Cooray A., 2008, ApJ, 689, 1
 Scannapieco C., Tissera P. B., White S. D. M., Springel V., 2006, MNRAS, 371, 1125
 Scherrer R. J., Bertschinger E., 1991, ApJ, 381, 349
 Shimasaku K., et al., 2006, PASJ, 58, 313
 Stark D. P., Loeb A., Ellis R. S., 2007, ApJ, 668, 627
 Storrie-Lombardi L. J., McMahon R. G., Irwin M. J., Hazard C., 1994, ApJ, 427, L13
 Trac H., Cen R., 2007, ApJ, 671, 1
 Trac H., Cen R., Loeb A., 2008, ApJ, 689, L81
 Yepes G., Kates R., Khokhlov A., Klypin A., 1997, MNRAS, 284, 235
 Zuo L., 1992, MNRAS, 258, 36
 Zuo L., Phinney E. S., 1993, ApJ, 418, 28